REMARKS

This Amendment is filed in response to the Office action dated July 22, 2008. All objections and rejections are respectfully traversed.

Claims 1-9 and 13-35 are in the case.

Claim 10-12 have been cancelled without prejudice

Request for Interview

The Applicant respectfully requests a telephonic interview with the Examiner after the Examiner has had an opportunity to consider this Amendment, but before the issuance of the next Office Action. The Applicant's undersigned may be reached at 617-951-2500.

Objection to the Specification

At paragraph 4 of the Office Action, Examiner objected to the specification as failing to provide antecedent basis for the claimed subject matter. Specifically, Examiner required correction to claims 33-34 for reciting "specific cluster partner before a higher virtual interface layer has..." Applicant respectfully traverses this objection. Applicant's specification states on page 4, lines 19-20,

> Thus, the cluster connection manager may be operational **before a higher level VI layer of the storage operating system** has fully initialized." Applicant further clarifies that "VI" is an abbreviation for "virtual interface", "These processes typically communicate with each other using a conventional Virtual Interface (VI). (Page 3, Lines 20-21)

Thus, it is the Applicant's belief that the specification provides proper antecedent basis for the claimed subjected matter.

Objection to the Drawings

At paragraph 5 of the Office Action, the Examiner objected to the Drawings under 37 U.S.C. 1.83(a) because it is the Examiner's belief that said drawings do not show every feature of the invention specified in the claims. Specifically, the Examiner points to claims 15-18, 20, 24, and 28 which claim "a predefined hardware address and predefined port number previously known to support a RDMA operations" and claims 33 and 34 which claim "a read operation directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and specific address that support a RDMA operations and performing a second…:" Applicant respectfully traverses this objection.

In Figure 2 of the Applicant's drawings, Applicant shows a schematic block diagram of an exemplary storage system. The storage system depicted contains a memory. One skilled in the art would understand a memory to have a plurality of possible "predefined hardware address." In addition the storage system depicted in Figure 2 also contains a Network Target Adaptor which would be known by those skilled in the art to contain a plurality of "predefined port numbers." Furthermore, Figure 4 specifically shows that operation 490 must be supported by RDMA operations. Applicant therefore believes the drawings to depict every feature of the invention specified in the claims.

Rejections Under §112 First Paragraph

At paragraph 6 of the Office Action, the Examiner rejected claims 33 and 34 under 35 U.S.C §112, first paragraph, as failing to comply with the written description requirement.

As noted above, Applicant respectfully traverses this rejection. Applicant's specification states on page 4, lines 19-20, "Thus, the cluster connection manager may be operational **before a higher level VI layer of the storage operating system** has fully initialized." Applicant further clarifies that "VI" is an abbreviation for "virtual interface" which is included in a previous paragraph stating, "These processes typically communicate with each other using a conventional Virtual Interface (VI)." (Page 3, Lines 20-21) Thus, it is the Applicant's belief that the specification provides proper antecedent basis for the claimed subjected matter.

Rejections Under §112 Second Paragraph

At paragraph 9 of the Office Action, the Examiner rejected claims 10, 11, and 12 under 35 U.S.C §112, second paragraph, as failing to comply with the written description requirement. Claims 10-12 have been cancelled without prejudice.

Rejections Under 35 U.S.C. §103(a)

At paragraph 19 of the Office Action, the Examiner rejected claims 1-3, 13-14, and 35 under 35 U.S.C. § 103(a) as being unpatentable over Craddock et. al, U.S. Publication No. 2003/0061296, published March 27, 2003, (hereinafter "Craddock") in view of Pandya et al., U.S. Publication No. 2004/0037319, published February 26, 2004, (hereinafter "Pandya").

Applicant's claimed novel invention, as set forth in representative claim 1, comprises in part:

> 1. A method for initiating a peer-to-peer communication session, the method comprising :
> ***attempting a first remote direct memory access (RDMA) read operation directed to a cluster partner having an operating system, the RDMA read operation bypassing the operating system;***

performing, in response to a successful first RDMA read operation, a first RDMA write operation to the cluster partner;

performing, in response to a successful RDMA write operation, a second RDMA read operation directed to the cluster partner; and

performing, in response to a successful second RDMA read operation, a second RDMA write operation to the cluster partner.

Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses I/O transactions which represent a unit of I/O work and typically contain multiple messages. Craddock provides a mechanism for initiating and completing one or more I/O transactions using memory semantic messages which are transmitted by a remote direct access (RDMA) operation. In particular, these I/O transactions **are read from a specific disk sector** onto a specific host memory location. Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector. Specifically, a process already running on a host first reserves a memory space for holding read data. The process then invokes a device driver associated with the storage device adapter, specifying that data from the storage device is to be read into read data memory space. At the close of the read transaction, the adapter generates a response and an **associated** write RDMA with an immediate work queue element that is interpreted and processed and transmitted via RDMA transfer to a host where it is stored in a location which was reserved for the response message.

Pandya merely discloses issuing an RDMA write command to a recipient. Then the recipient executing a write command by allocating an RDMA buffer to be received by the write and then requesting a RDMA read.

Applicant respectfully urges that neither Craddock nor Pandya show Applicant's claimed novel use of ***attempting a first remote direct memory access (RDMA) read operation directed to a cluster partner having an operating system, the RDMA read operation bypassing the operating system.***

Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) In Applicant's system, ***the RDMA read operation bypasses the operating system.*** That is, Applicants invention uses specific ports and network addresses so that it knows that the connection will be reliable based on the protocols known to support RDMA operations. That is, it does not rely on the Virtual Interface Layer to make the determination. Specifically, ***the method initiates a peer-to-peer communication session by first attempting a remote direct memory access read operation directed to a predefined hardware address and a predefined port number which is known to support remote direct access memory operations. If a peer-to-peer connection is successful, then a RDMA write operation is performed wherein it too is directed to the predefined hardware address and the predefined port number.***

Craddock does not disclose accessing in any data from any location while at the same time bypassing the operating system. Craddock transfers messages via the verbs interface and message and data service. Craddock states in paragraph 0050 that this system uses both a verbs interface layer and a message and data interface layer to operate. These layers in combination act like virtual interface layer because they define the destinations (i.e., addresses and ports) which must be used in order for the RDMA operation to function properly. Applicant's invention accesses the cluster partner and bypasses the operating system therefore it does not use the virtual interface layer to operate. Additionally, Pandya merely discloses issuing a second RDMA write operation.

Applicant respectfully urges that Craddock and Pandya either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed ***attempting a first remote direct memory access (RDMA) read opera-***

tion directed to a cluster partner having an operating system, the RDMA read operation bypassing the operating system.

At paragraph 12 of the Office Action, the Examiner rejected claims 15-19 under 35 U.S.C. § 103(a) as being unpatentable over Craddock in view of Prakash et al., U.S. Patent No. 6,434,626, issued on August 13, 2002 (hereinafter "Prakash").

Applicant's claimed novel invention, as set forth in representative claim 15, comprises in part:

> 15. A method comprising :
>
> ***initiating a peer-to-peer communication session which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation.; and***
>
> performing, in response to a successful step (a), a first remote direct memory access write operation directed to the predefined hardware address and the predefined port number.

As discussed above, Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses …memory semantic messages which are … **read from a specific disk sector…**. Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector.

Prakash merely discloses initiating a peer-to-peer communication session.

Applicant respectfully urges that neither Craddock nor Prakash show Applicant's claimed novel use of ***initiating a peer-to-peer communication session which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation.***.

As noted above, Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) ***which bypass the storage system's operating system.***

Craddock does not disclose accessing any data from any location ***which bypass the storage system's operating system.*** Craddock transfers messages via the verbs interface and a message and data interface layer service. Craddock states in paragraph 0050 that this system uses both a verbs interface layer and a message and data interface layer to operate. These layers in combination act like virtual interface layer because they define the destinations (i.e., addresses and ports) which must be used in order for the RDMA operation to function properly. Applicant's invention accesses the cluster partner ***which bypasses the storage system's operating system*** (i.e., where the virtual interface layer runs). Furthermore, Prakash, merely initiating a peer to peer communication session and does not discusses bypassing the storage systems operating system.

Applicant respectfully urges that Craddock and Pandya either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed ***initiating a peer-to-peer communication session which bypasses an***

operating system on a storage system ***by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation.***

At paragraph 32 of the Office Action, the Examiner rejected claims 10-12 and 20-31 under 35 U.S.C. § 103(a) as being unpatentable over Craddock and Prakash in view of Sutherland et al., U.S. Publication No. 2002/0114341, published on August 22, 2002 (hereinafter "Sutherland"). Claims 10-12 have been cancelled without prejudice. Applicant has amended the remaining claims and believes them to be in condition for allowance.

Applicant's claimed novel invention, as set forth in representative claim 20, comprises in part:

> 20. A system configured to establish reliable peer-to-peer communication among storage systems of a clustered environment, the system comprising:
>
> a peer process executing on each storage system partner having an operating system; and
>
> ***a cluster connection manager executing on each storage system partner, the cluster connection manager establishing a reliable peer-to-peer connection between each peer process by connecting to a predetermined port number using a predetermined network address, the reliable peer-to-peer connection bypassing the operating system.***

As discussed above, Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses…memory semantic messages which are…**read from a specific disk sector**. Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector.

Pandya merely discloses issuing an RDMA write command to a recipient. Then the recipient executes a write command by first allocating an RDMA buffer to receive the write and then requesting a RDMA read. Again, however, nowhere in the Pandya disclosure is there a suggestion that any of these processes bypass the operating system of the storage system.

Sutherland merely discloses a cluster connection manager.

As argued above neither Craddock nor Pandya disclose bypassing the storage systems operating system to access a predetermined port. Sutherland also does not discloses this as Sutherland merely discloses the use of a cluster connection manager. Sutherlands' cluster connection manager however does not execute on each storage system partner to establishes a reliable peer-to-peer connection between each peer process by connecting to a predetermined port number using a predetermined network address, ***the reliable peer-to-peer connection bypassing the storage system's operating system***.

Applicant respectfully urges that neither Craddock, Prakash, nor Sutherland either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed use of ***a cluster connection manager executing on each storage system partner, the cluster connection manager establishing a reliable peer-to-peer connection between each peer process by connecting to a predetermined port number using a predetermined network address, the reliable peer-to-peer connection bypassing the operating system.***

At paragraph 42 of the Office Action, the Examiner rejected claim 33 under 35 U.S.C. § 103(a) as being unpatentable over Craddock and Pandya in view of Boyd et al., U.S. Publication No. 2004/0049600, published on March 11, 2004 (hereinafter "Boyd").

Applicant's claimed novel invention, as set forth in representative claim 33, comprises in part:

> 33. A method comprising:
>
> ***initializing a first remote direct memory access (RDMA) read operation that bypasses the operation system and is directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and a specific address that support a RDMA operations; and***
>
> performing a second RDMA read operation directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and a specific address that support a RDMA operations.

As discussed above, Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses ...memory semantic messages which are...read from a specific disk sector.... Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector.

Pandya merely discloses issuing an RDMA write command to a recipient. Then the recipient executes a write command by first allocating an RDMA buffer to receive the write and then requesting a RDMA read.

Boyd discloses a memory management system which uses a virtual interface layer to assign virtual addresses to physical addresses via a memory table. The memory management system is connected to an IP net (Fig 1). The IP net allows the user to bypass the host processor node and directly access the RAID subsystem. Within the RAID subsystem is a memory and a processor, the memory contains an operating system which all accesses to the RAID system must still go through. That is, although the user can bypass the host processor's operating system, the user may still not bypass the operating system of the RAID subsystem.

Applicant respectfully urges that neither Craddock, Pandya, nor Boyd show Applicant's claimed novel use of ***initializing a first remote direct memory access (RDMA) read operation that bypasses the operation system and is directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and a specific address that support a RDMA operations.***

As noted above, Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) ***which bypasses the operating system on the storage system.***

Boyd merely discloses what a virtual interface layer does. Applicant is not claiming a virtual interface layer but rather a method for bypassing the operation system on a storage system before the virtual interface layer has initialized. Boyd allows the user to bypass the host processors operating system but as can be seen by Boyd's Figure 1, the user may not bypass the storage systems operating system. The user can only directly access an IPSOE (Figure 1, 172).

Applicant respectfully urges that neither Craddock, Prakash, nor Boyd either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed use ***initializing a first remote direct memory access (RDMA) read operation that bypasses the operation system and is directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and a specific address that support a RDMA operations.***

Conclusion

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Stephen D. LeBarron/

Stephen D. LeBarron
Reg. No. 62,479
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
Telephone: (617) 951-2500
Facsimile: (617) 951-3927